Remarks

Reconsideration and allowance of pending claims 1-25 are respectfully requested.

In the Office Action, claims 1-6, 11-13 and 18-20 were rejected under 35 U.S.C. §103(a) as being unpatentable over Torbjørnsen et al. (U.S. Patent No. 5,555,404; hereinafter Torbjørnsen) in view of Biliris et al. (U.S. Patent No. 5,966,706; hereinafter Biliris), while claims 7-10, 14-17 and 21-24 were rejected under 35 U.S.C. §103(a) as being unpatentable over Torbjørnsen in view of Biliris and further in view of Badovinatz et al. (U.S. Patent No. 5,805,786; hereinafter Badovinatz), and claim 25 was rejected under 35 U.S.C. §103(a) as being unpatentable over Torbjørnsen, in view of Biliris and further in view of Duprey et al. (U.S. Patent No. 6,671,705; hereinafter Duprey). Each of these rejections is respectfully traversed, and reconsideration thereof is requested.

As recited in claim 1, for example, Applicants' invention comprises a technique for recovering from failures within a shared nothing distributed computing environment. The technique includes detecting multiple overlapping failures within the shared nothing distributed computing environment, and automatically recovering from the multiple overlapping failures. Thus, Applicants' technique is a multiple fault tolerant approach for a shared nothing distributed computing environment. Further, Applicants' approach specifies automatic recovery from the failures wherein one or more transactions affected by the failures are automatically executed to completion without rolling back the transactions and without requiring a reposting of the transactions. Applicants respectfully submit that a careful reading of the applied art fails to uncover any teaching or suggestion of these concepts within the context of a multiple fault tolerant approach for a shared nothing distributed computing environment.

An "obviousness" determination requires an evaluation of whether the prior art, taken as a whole, would suggest the claimed invention, taken as a whole, to one of ordinary skill in the art. In evaluating claimed subject matter as a whole, the Federal Circuit has expressly mandated that functional claim language be considered in evaluating a claim relative to the prior art. Applicants respectfully submit that the application of these standard to the independent claims

presented herewith leads to the conclusion that the recited subject matter would not have been obvious to one of ordinary skill in the art based on the applied patents.

For example, a careful reading of Torbjørnsen fails to uncover any discussion of a failure recovery approach that is applicable to multiple overlapping failures (let alone a multiple overlapping failure recovery approach wherein one or more transactions affected by one or more failures of the multiple overlapping failures are automatically executed to completion without rolling back the one or more transactions and without requiring a reposting of the one or more transactions). Torbjørnsen expressly teach otherwise at Column 5, lines 47-55, where it is stated that the corrective online repair approach disclosed therein is single fault tolerant. As noted, if a second fault were to occur while the repair of the first fault is taking place, the system is vulnerable to failure. This is in contrast to Applicants' invention which is directed to a high reliability system wherein multiple overlapping failures may occur, be detected and automatically recovered from in such a way that one or more transactions affected by the failures are automatically executed to completion without rolling back the transactions and without requiring reposting of the transactions.

Biliris describes a distributed database management computer system which includes a plurality of nodes and a plurality of database pages. When a first node in the computer system updates a first database page, the first node generates a log record. The first node determines whether it manages the first database page. If the first node determines that it manages the first database page, the first node writes the log record to a log storage local to the first node. However, if the first node determines that it does not manage the first database page, the first node then determines whether it includes a local log storage. If the first node includes a local log storage, the first node writes the log record to the local log storage, even if the first node does not manage the first database page. If the first node does not include a local log storage, the first node sends the log record to a second node managing the first database page. (See abstract.)

The Office Action acknowledges that Torbjørnsen does not teach recovery from multiple overlapping failures. Biliris is cited for allegedly teaching recovery from multiple overlapping failures at Col. 12, line 58 through Col. 13, line 46. This characterization of the teachings of Biliris, to the extent deemed applicable to the claimed invention, is respectfully traversed. The cited lines of Biliris describe a multiple node crash recovery process wherein nodes continue to access pages in local cache while the remaining nodes are in the process of recovery. In Biliris, data is unavailable until a failing node is brought back online, examines its log, and then completes a reposting operation.

Applicants respectfully submit that there is no teaching or suggestion in Biliris of automatic recovering from multiple overlapping failures, wherein one or more transactions affected by one or more failures of the multiple overlapping failures are automatically executed to completion without rolling back the one or more transactions and without requiring a reposting of the one or more transactions. Biliris teaches at Col. 13, lines 28+ that:

> Although each crashed node lost its DPT during the crash, a superset of each node's DPT can be reconstructed by scanning the node's log file. In particular, each crashed node scans its log by starting from the last complete checkpoint and updates the DPT stored in that checkpoint by inserting new entries for the pages that do not have an entry and are referenced by the examined log records.

Applicants respectfully submit that these lines expressly teach the reposting of information to the log, and thus teach away from Applicants' functionality wherein transactions affected by one or more failures of the multiple overlapping failures are automatically executed to completion without rolling back the one or more transactions and without requiring a reposting of the one or more transactions. Since this functionality recited by Applicants in their independent claims is not taught or suggested by Torbjørnsen or Biliris in a multiple overlapping failure environment, Applicants respectfully submit that their invention would not have been obvious to one ordinary skill in the art based thereon. There is no suggestion in Torbjørnsen or Biliris, individually or in combination, of a technique for recovering from multiple overlapping failures wherein one or more transactions affected by one or more failures of the multiple overlapping failures are automatically executed to completion without rolling back the one or more transactions and without requiring a reposting of the one or more transactions.

Additionally, Applicants respectfully submit that there is no suggestion or implication in Torbjørnsen or Biliris how one of ordinary skill in the art would modify the Torbjørnsen system to make the approach described therein a multiple fault tolerant approach as suggested in the Office Action. Torbjørnsen expressly teaches that the mirroring approach described therein is a single fault tolerant approach meaning that only a single fault at a time can be addressed. This is contrasted with Applicants' invention wherein multiple overlapping failures within the shared nothing distributed computing environment can be detected and recovered from automatically, wherein one or more transactions affected by one or more failures of the multiple overlapping failures are automatically executed to completion without rolling back the one or more transactions and without requiring a reposting of the one or more transactions. Neither Badovinatz nor Duprey is cited in the Office Action for this aspect of Applicants' invention, and neither is believed applicable thereto.

Badovinatz is cited in the Office Action for teaching recovery in a distributed computing environment, wherein the recovery includes electing a coordinator from among at least one surviving member, while Duprey is cited for a remote mirroring system, which includes a failure recovery technique for a distributed synchronous transaction system. Without acquiescing to the above characterizations of the teachings of these two patents, Applicants respectfully submit that neither patent addresses an approach for detecting multiple overlapping failures and automatically recovering from the multiple overlapping failures within a shared nothing distributed computing environment, let alone automatically recovering from those failures so that transactions affected by the failures are automatically executed to completion without rolling back the one or more transactions and without requiring a reposting of the one or more transaction. The protocol necessary to achieve this automatic failure recovery in the presence of multiple overlapping failures would not have been obvious to one of ordinary skill in the art based upon the art of record.

For the above reasons, Applicants respectfully submit that the independent claims presented are patentable over the applied art. The dependent claims are allowable for the same reasons as the independent claims, as well as for their own additional characterizations.

The application is believed to be in condition for allowance, and such action is respectfully requested. Applicants' undersigned attorney is available should any remaining issue require resolution.

Respectfully submitted,

Kevin P. Radigan

Kevin P. Radigan
Attorney for Applicants
Registration No. 31,789

Dated: August 31, 2004

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203
Telephone: (518) 452-5600
Facsimile: (518) 452-5579